AGRICULTURAL ECONOMICS

Agricultural Economics 43 (2012) supplement 73–84

Social learning and technology adoption: the case of coffee pruning in Peru

Jeremy G. Weber*

U.S. Department of Agriculture, Economic Research Service, 1400 Independence Ave. SW, Washington, DC 20250-1800, USA

Received 30 January 2012; received in revised form 9 May 2012; accepted 12 June 2012

Abstract

A unique natural experiment involving a coffee pruning technology is used to study social learning. The yield effects of pruning take two years to appear, a characteristic that aids in identifying social learning apart from correlated unobservable variables that are a concern in the social learning and technology adoption literature. Panel data are employed that start with a private initiative which introduced systematic pruning in central Peru and that contain the population of participating growers. Results show a jump of at least 0.15 in the probability of adoption two years after the first pruning in a grower's group.

JEL classifications: D83; Q12

Keywords: Social learning; Technology adoption; Coffee, Peru

1. Introduction

Studies have explored how learning from others can affect choices about technology, especially in agriculture. Besley and Case (1994) and Foster and Rosenzweig (1995) study the adoption of high yielding seed varieties in India. For both, social learning occurs when the behavior of neighbors (e.g., area planted in a new seed variety) informs a farmer's belief about the profitability of new seeds, which contrasts with individual learning where a farmer's past experience informs his current decisions. Munshi (2004) shows how farmers may learn less from their neighbors if the performance of the technology depends on unobserved characteristics of the neighbors. Bandiera and Rasul (2006) and Maertens (2009) add to the literature by relating the behavior of farmers with the behavior of peers that the farmers themselves identify. More recently, Conley and Udry (2010) find that farmers change their fertilizer use in response to the input levels used by information neighbors who achieved better than expected results in previous periods.

*Corresponding author. Tel.: 202-694-5584; fax: 202-694-5774. *E-mail address*: jeweber@ers.usda.gov (J. G. Weber).

Data Appendix Available Online

A data appendix to replicate main results is available in the online version of this article. Please note: Wiley-Blackwell is not responsible for the content or functionality of any supporting information supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.

Identifying social learning, however, is challenging. As Manski (1993) highlights, an individual's behavior may be correlated with the average behavior of the group because individuals face a similar environment (contextual effects) or have similar exogenous characteristics (correlated effects). Furthermore, serially correlated unobservable variables could induce correlation between an individual's current behavior and the past behavior of neighbors. Finding a correlation between group adoption and individual adoption, contemporaneous or not, may therefore have little to do with social learning.

This article uses a natural experiment that provides a unique opportunity for identifying social learning. The coffee trading company, Volcafe, and the Italian coffee roaster, Lavazza, financed a project to work with small scale growers in central Peru.[1] The project's technical advisor came from Costa Rica and introduced a pruning practice not used in Peru at the time. The dynamic and observable effects of pruning make for a fruitful case study of social learning—seeing a recently pruned plant provides no direct information about the effects of pruning on yields; seeing a plant pruned two years ago is in most cases conclusive. The revelation of information implies a discrete increase in adoption at a certain time if social learning is important. The jump in adoption, not expected in the absence of social learning, is consistent with the finding of Brock and

[1] For an overview of the project, visit http://www.lavazza.com/corporate/en/company/socialreponsability/Il_progetto/

 DOI: 10.1111/j.1574-0862.2012.00621.x

Durlauf (2010) that under relatively weak assumptions social interactions can produce discontinuities in adoption curves that would not occur in their absence. Furthermore, unless a grower is the first to prune in his peer group, individual learning provides little incentive to prune since a grower's own recently pruned plants provide far less information than a neighbor's plants that have fully recovered—a characteristic of pruning that aids in identifying social learning apart from individual learning.

My empirical analysis uses a unique panel data set. While only three geographically isolated groups are observed, the data set includes the population of growers that participated in the project and covers six years starting when project extension agents first introduced systematic pruning. Furthermore, the rollout of the project from one group to another provides an exogenous variation in the timing of the first prune in the group, which aids in separating the effect of revealed information from the first pruned coffee plants from year-specific shocks. All estimation approaches based on the timing of the first prune in a group reveal a clear pattern: there is a decrease in adoption from the initial year, when at least one person in the group prunes (Year 0), to the following year (Year 1). But in the following year (Year 2), a discrete increase in adoption occurs, which is not seen from Year 2 to Year 3. In contrast, an alternative model of social learning based on accumulated group experience, which is popular in the literature but ignores the delayed recovery of pruned plants, fails to provide evidence of social learning.

2. Pruning in central Peru: A natural experiment for studying social learning

Lavazza and Volcafe financed a project to help growers in the department of Junin, Peru improve the sustainability of their coffee farms. I refer to this privately funded sustainability initiative as 'the project'. Growers initially learned about the project through public meetings that outlined activities and goals. The project started working in one district in 2005 and then expanded to two more districts in 2006 and 2007 (the first three largest territorial divisions in Peru are department, province, and district). In total the project, which formally ended in 2009, had 315 participating growers distributed across the three districts.

The project had two core components: improve the social and environmental conditions of coffee growing households using Rainforest Alliance's Sustainable Coffee program and increase the profitability of their farms through better practices. The key strategy for increasing profitability was to increase yields through systematic pruning. Coffee plants reach peak production around five years of age after which yields can decline, though how much depends on the management of the plant (fertilizer use, pest management, etc.). According to research by the Colombian National Center for Coffee Research, pruning can cause older plants (at least five years old) to produce quantities on par with plants in their first years of production, thus delaying the decline in yields associated with aging plants. Research recommends pruning trees once every three to five years until they reach 20 years, after which growers should replace them with new plants.[2]

Systematic pruning involves annually cutting a share of coffee plants to about a meter in height; for example, cutting all plants in every third row. The recommended time to prune is immediately following the harvest. A plant pruned after the harvest in Year 0 produces nothing in Year 1 but returns to production in Year 2. Aside from inferences based on the actions of others, nothing can be learned by observing a recently pruned plant. In Year 1, the plant starts to recover, with an increase in its foliage, which indicates that pruning does not kill the plant but little else. In Year 2, the plant produces much more than it did prior to pruning. Even without observing yields, the thick foliage of a plant pruned two years ago is a strong signal that pruning improves plant vitality.

While variations of systematic pruning are common in Costa Rica and Colombia, two countries with well developed coffee sectors, the practice was not used in Peru prior to 2005. There, the conventional method has been selective pruning, where a grower annually selects and cuts a few branches of each plant, but does not cut it down to a much smaller size. Growers initially viewed systematic pruning (here after 'pruning') with skepticism,[3] in part because the practice was introduced from the outside by a coffee expert from Costa Rica. In time, skepticism has abated as evidenced by the cumulative adoption rate in Fig. 1, and the practice has become part of a national plan to improve yields (Junta Nacional del Café—Peru, 2011; Agrobanco, 2011).

The crop and context make for a policy-relevant study of social learning, especially from an international rural development perspective. Coffee is a perennial crop that requires substantial upfront investment for returns that take several years to materialize. And although coffee growing households in Latin America are generally not the poorest of the rural poor, the project targeted lower-income smallholders (three or four hectares of coffee in addition to small plots of bananas and avocados), who would be particularly risk adverse, especially concerning management of their main cash-generating asset, coffee plants. Consequently, finding that information revealed by neighbors' actions can cause growers to start cutting their plants in half would attest to the power of seeing neighbors succeed with a technology.

I argue that the exogenous introduction of the technology combined with its distinct pattern of revelation of information provides a unique opportunity to identify social learning under

[2] See recommendations by Centro Nacional de Investigaciones de Café, www.cenicafe.org, in their page "Sistemas de Produccion" and Avance Tecnico 0215.

[3] One grower commented that initially he preferred to hire someone to prune his trees since he could not bring himself to cut them. Another grower's wife threatened to divorce him if he pruned their plants.

a weak assumption about correlated unobservable variables. Much of the social learning and technology adoption literature struggles to separate the effect of average group choices on individual behavior (an endogenous social effect) from the effect of unobservable variables operating at the group level. As Conley and Udry (2010) state, "correlated unobservables are a general problem in the literature on agrarian technology."

Two empirical approaches dominate the literature. One approach relates lagged village aggregates like acreage planted or average yields to the adoption decisions of an individual farmer in the village (Foster and Rosenzweig, 1995; Munshi, 2004; Moser and Barrett, 2006). A challenge of this literature is dealing with serially correlated unobservable variables that would lead to correlation between lagged aggregate measures and individual behavior. The other approach identifies an individual's reference group and links information from the reference group to the individual's behavior (Bandiera and Rasul, 2006; Maertens, 2009; Conley and Udry, 2010).

My empirical strategy is slightly different from the literature that links past group behavior to individual decisions. What matters here is when someone in the group started pruning, not the group's accumulated experience. As suggested above, if social learning is important, the initial pruning in the group will have a positive effect on adoption two years later, but little or no effect in the intervening year. This approach to identifying social learning rests on assuming that correlated unobservable variables did not induce this unique adoption pattern. To illustrate the potential issue with correlated unobservable variables, Conley and Udry posit an example where neighboring farmers i and j each use a low level of fertilizer at time t. At t+3, farmer j experiences a positive weather shock and uses a high level of fertilizer. Because of positive spatial and serial dependence, farmer i is also likely to experience a positive shock at, say, $t + 8$, and use a higher level of fertilizer. Thus, what looks like farmer i learning from farmer j's innovation is really a product of a correlated unobservable variable. For pruning, the shock that contributes to the initial adoption in the group would have to disappear in the following year but then reappear to some degree two years later when growers can observe the effects of pruning.

To be conceptually clear on identification, social learning is a complex phenomenon that researchers outside of economics have studied in detail for some time (Miller and Dollard, 1941). Recent economic work recognizes many aspects of the diffusion of private information, perhaps involving strategic interactions and multiple signals (words, actions, effects of actions) (Chamely, 2004). I claim to identify only one form of learning where a grower learns about the profitability of a new technology by observing the fruit of his neighbors' prior actions. A grower may come to knowledge of pruning through conversation, local newspapers, or town meetings, but only first-hand observation of the effects of pruning are likely to induce adoption. Claims of the effectiveness of pruning become credible when verifiable evidence—a pruned plant with thick foliage or a bountiful harvest—support them.

3. Conceptual framework

Project extension agents recommended that growers annually prune a third of their farm. While there was no formal incentive (or disincentive) to follow the recommendation, growers seemed to have interpreted their options as pruning a considerable section each year (i.e., more than a few plants) or not pruning. I keep the conceptual model general by treating pruning as a continuous variable—the share of the farm that a grower prunes. The decision involves dynamic considerations since there is lag between adoption decisions and outcomes. Pruning occurs after the harvest and therefore does not affect output in that year. Furthermore, the perennial nature of coffee plants implies that pruning just once will affect production well into the future.

I limit the time horizon so that the grower does not prune the same plant twice from t = 0...T. Thus, I assume that at some t > T, the grower decides to prune the first section again or replace it with new plants. The variable a_{it} is the share of plants that a grower prunes at time t and ranges from 0 to 1, inclusive. Output is given by the production function $F_i(.)$ and is a function of the area pruned in each previous period $a_{i\tau=0\ldots t-1}$, inputs X_{it}, and a random variable μ_{it}. (Upper case letters represent vectors; lower case letters represent scalars). The production function is grower specific and captures soil quality, elevation, and other relatively fixed factors that determine the productive capacity of a stand of plants. The random variable (μ_{it}) reflects uncertainty over the effect of pruning on output and is the object of learning. Grower i at time $0 \leq t \leq T$ earns a profit Π_{it} of

$$\Pi_{it} = p_{ct} F_i\,(a_{i\tau=0\ldots t-1}, X_{it}, \mu_{it}) - P_{xt} X_{it}, \tag{1}$$

where p_{ct} is the price of coffee and P_{xt}, is a vector of input prices.

Assuming a time separable utility function U that allows for general risk preferences, the grower maximizes the present value of utility from a discrete series of profits over a finite horizon. In each period the grower uses the available information to decide what to do; E_t represents the grower's subjective expectations at time t and can be informed by his own actions (individual learning) or by information from others (social learning). With δ representing the discount rate, the objective function at time t is

$$\sum_{t=0}^{T} \delta^t E_t\,(U\,(\Pi_{it}\,(.))). \tag{2}$$

Because pruned plants yield nothing in the first year and rural credit markets are imperfect, the grower faces a borrowing constraint $\Pi_{it}(a_{it-1}) > \beta(W_{it})$ where W_{it} reflects a household's liquidity needs and wealth.

The time-frame T is chosen to be small enough that the grower would not want to prune the same plants twice prior to T. Thus, if a grower has not pruned before t, future pruning decisions are restricted such that

$$a_{it+1} \leq 1 - a_{it}. \tag{3}$$

The restriction in (3) simply means that if a grower started pruning in t by pruning a third of his plants, then in the following year he will at most prune the remaining two-thirds.

Growers in the same group initially receive the same information about pruning from project extension agents, though individual growers may weigh the information differently which could explain why not all growers start pruning at the same time. The information provided from individual experiments with pruning could motivate the first adopters. My focus, however, is not when pruning started in a group, but the adoption pattern that followed.

Given an initial assessment of information from extension agents, the grower's expectation about the random variable μ_{it} will evolve as information accumulates through individual or social learning. Because a plant pruned earlier provides more information than a plant pruned more recently, the years since a plant was pruned shapes the grower's information set. To allow for different effects of observing one's own pruned plants and observing a neighbor's plants, let y_{it} be the years since grower i started pruning and $\widetilde{y}_{it}$ be the years since the first prune in grower i's group. The y variables have an initial value of zero and evolve in the following manner:

$$y_{it+1} = \begin{cases} y_{it} + 1 & \text{if } a_{i\tau} > 0 \text{ for some } \tau \leq t \\ 0 & \text{otherwise} \end{cases} \tag{4a}$$

$$\widetilde{y}_{it+1} = \begin{cases} \widetilde{y}_{it} + 1 & \text{if } a_{j\tau} > 0 \text{ for some } j \in group(i) \\ & \quad \text{and } \tau \leq \text{t} \\ 0 & \text{otherwise} \end{cases} \quad . \tag{4b}$$

Let I be a generic function relating variables reflecting the age of oldest prune to the random variable μ_{it} that reflects uncertainty about pruning.

$$\mu_{it} = I(y_{it}, \widetilde{y}_{it}). \tag{5}$$

In the year of the first prune in the group, the value of $I(y_{it} = 0, \widetilde{y}_{it} = 0)$ reflects the grower's assessment of information from extension agents and any inferences from observing the decisions of neighbors. In the following year, relatively little can be learned about the effects of pruning by observing a plant pruned a year ago. In a Bayesian framework, the grower would update his prior belief informed by extension agents and inferences from observing neighbors with the signal sent by observing a one-year prune. A plant pruned two years ago, however, provides credible information on the effect of pruning.

Suppressing the notation for prices, I incorporate Eqs. (1) to (5) into Bellman's equation that expresses the value function Z at time t as a function of y_{it} and $\widetilde{y}_{it}$.

$$Z_{it}(y_{it}, \widetilde{y}_{it}) = \max_{Xit,ait} E_t \begin{pmatrix} U\left(\Pi_{it}\left(a_{i\tau=0\ldots t-1}, X_{it}, y_{it}, \widetilde{y}_{it}\right)\right) + \\ \delta Z_{it+1}\left(g(a_{i\tau=0\ldots t}, X_{it}, y_{it}, \widetilde{y}_{it}\right)\right) \end{pmatrix}$$
$$s.t. \Pi_{it}(a_{it-1}) > \beta(W_{it}). \tag{6}$$

The state variables are the area previously pruned, which specifies the area still available to be pruned as shown by Eq. (3), and the information available based on the oldest prune of grower i and grower i's group, $I(y_{it}, \widetilde{y}_{it})$. The control variables chosen in the current period are the area to be pruned (a_{it}) and input use (X_{it}). The generic function g relates the state and control variables in the current period to next period's value function.

If no one in the group has pruned before, the grower's current pruning decision will affect future pruning decisions by increasing knowledge of μ, which is implicit in g. If the grower has pruned before, knowledge about μ increases mechanically as previously pruned plants recover (captured by y_{it}). If someone in the group has pruned previously, knowledge about μ increases in a similar manner (captured by $\widetilde{y}_{it}$).

The grower will prune in the current period if next period's value function from pruning exceeds that from not pruning:

$$V_{it+1}(a_{it}, .) - V_{it+1}(0, .) > 0, \tag{7}$$

where $V_{it+1}(a_{it}, .)$ is the conditional value function defined as

$$V_{it+1}(a_{it}, .) = \max_{Xit+1} E_{t+1}\left(U\left(\Pi_{it+1}(.)\right) + \delta Z_{it+2}(.)\right). \tag{8}$$

4. Data and descriptive statistics

The project started in two communities in district 1 in 2005 and then expanded within the district and to another district in 2006. In 2007, it expanded to another district. At the smallest geographical level, the project worked with local growers associations, and while it is possible to use the association as a grower's relevant reference group, responses to a 2010 survey question concerning visits to early adopters and conversations with key informants revealed that some growers visited farms of growers from nearby associations to observe the effects of pruning. Unmeasured spillovers of information from one district into another could be captured by treating the entire project as one group and testing if the temporal pattern of adoption was consistent with the distinct pattern of revelation of information from a recovering pruned plant. Doing so, however, would preclude separating discontinuities in temporal adoption patterns introduced by social learning from year-specific shocks affecting adoption. To limit information spillovers but to have temporal variation in the start of the first prune in a group, I organize growers by district, which are geographically isolated from each other given the region's few roads. The main town of each district is about an hour and a half drive (in good conditions) from the main town of the other two districts.

I use information collected by project extension agents over the course of the project and from a survey conducted in 2010 by hired enumerators. Information collected by the project includes a registry of pruned area for each grower, yield information, and a baseline survey with basic household information like the number and age of household members and a breakdown of land holdings. There are good reasons to believe that the extension agents collected reliable information. The agents were in some cases coffee growers themselves (though not project participants) and were familiar with the agronomy, local terminology, and basic economics of coffee growing in the region. Perhaps more importantly, agents collected the baseline data knowing that external inspectors (e.g., for the Rainforest Alliance certification) would check some of the information. For example, extension agents had to walk the perimeter of a grower's farm and draw a croquis indicating land use (coffee, forest, other crop land) and area (number of hectares). The inspector would later verify the information on randomly selected farms.

I coordinated the 2010 survey using hired enumerators, some of whom had previously worked for the project. The survey occurred after the harvest in 2010 (August and September) and asked growers about the 2010 season in detail and for basic production and pruning information from 2006 to 2009 in addition to socio-economic variables and information on visits to the parcels of early adopters. Enumerators surveyed 236 of the 315 growers who had initially participated in the project. While some initial participants had died (4), refused to give information (7) or had sold their land (11), the main reason initial participants were not surveyed was because the grower was not at the farm at the time of the visit—after the harvest many growers leave their farms to live and work elsewhere. Baseline information for surveyed and nonsurveyed growers reveals statistically similar group means for ten farm and household variables. (Barham and Weber (2012) use the same survey data and make the comparison in Appendix B of their article).

Information on pruning prior to 2009 comes from the pruning registry maintained by the project. Because it is not recall data and was carefully updated by project extension agents based on repeated field visits, the pruning registry information is the most reliable source of pruning information for growers while they participated in the project. However, 116 growers left the project after joining, creating gaps in the pruning registry. More than half (65) were surveyed in 2010, and I use the responses to fill the registry's gaps.

Cases with pruning information from the registry and the 2010 survey permit assessing the reliability of recall data on pruning. I calculate an indicator for whether a grower pruned in a given year according to the survey and registry and measure reliability as the percent of cases where the two indicators agree. Reliability is high in 2006 and 2007 at 91 and 80% but it decreases to 56% in 2008. Though confusing at first glance, the declining reliability stems from greater pruning in 2007 and 2008 and the limited ability of growers to recall that they had pruned in earlier years. In 2008, more than four-fifths of the conflicting cases involved the survey indicating that the grower had not pruned but the registry indicating that they did. Possible underreporting in 2009, when only information from the 2010 survey is used, would cause a downward bias in estimation of the social learning effect if the majority of growers gained access to a two-year-old prune in 2009. But by 2009, 93% of growers already had access to a two-year-old prune. Thus, measurement error from recall data may soften the increase in adoption for the later years but have little effect on the first three years, which I am relying on for identifying social learning. It is important to note that the key explanatory variable for the empirical analysis, the timing of the first prune in the group, is always based on nonrecall data.

Table 1
A description of the sample ($N = 315$)

	Median	Mean	SD	Units
Coffee Land	3.0	4.2	4.2	Hectares
Forest Land	0.5	2.8	5.2	Hectares
Other Crop Land	0.0	0.7	1.7	Hectares
Total Land	5.0	7.7	7.4	Hectares
Gross Income – Other Crops[a]	340	1063	4822	US$
Coffee Yield[b]	400	588	955	Kg/Ha
Elevation	1.4	1.4	0.1	Thousands of Meters
Labor (Hhld members ages 10–65)	3.0	3.3	2.0	Number
Age of Household Head[b]	43.0	45.2	12.9	Number

Source: project baseline survey; tabulations by author.
[a]Households often have banana or avocado growing among coffee plants, which is why median *Gross Income – Other Crops* is positive but median *Other Crop Land* is zero.
[b]*Yield* is missing for two observations. Six observations lack the age of the household head.

Concerning the growers who left and were not surveyed in 2010, for 2006 to 2010 the number of grower-year observations where it is unknown if the grower has ever pruned is 4, 21, 23, 49 and 49. I retain these growers and assume that they did not prune. In the empirical section, I assess the robustness of the results to this assumption.

Table 1 presents descriptive statistics for the variables from the baseline survey conducted in most cases following the 2005 season. The sample consists of all 315 growers who initially joined the project.[4] The baseline for Group 3 occurred in January of 2007 and corresponds to the 2006 growing year.

Participant farms range from 980 to 1750 meters in elevation. Coffee is the mainstay of the economy in this agro-ecological niche, and Junin consistently produces more coffee than any other region in Peru. Unsurprisingly, coffee dominates household land portfolios—median coffee land is 60 of median total land holdings. While coffee is the main source of income, many households depend on banana and avocado for a 'caja chica,' a small but steady source of cash. Of the 228 farms that reported noncoffee crop income, 197 had income from selling bananas. Both banana and avocado are perennial crops that take time to reach the production stage but then consistently yield fruit.

[4] Two grower associations had a commercial relationship with the cooperative formed by the project but are not included given their lack of participation in project activities.

Table 2
A description of the groups

	Growers	Year project participation began	Year of access to 2 year old prune	Cumulative adoption rate by 2010
Group 1	199	2005	2007	50%
Group 2	21	2006	2009	29%
Group 3	95	2007	2008[a]	58%

Source: Project records, pruning registry, and 2010 survey; tabulations by author.
[a]The NGO implementing the project coordinated and financed a formal visit of growers from Group 3 to early adopters in Group 1 after the pruning season in 2007. See text for details.

Because coffee pruning removes plants from production for a year, one might expect household income diversification to influence adoption decisions.

Fort and Ruben (2009) sampled growers in the same area to study the impact of Fair Trade. Their sample included Fair Trade, organic, and conventional growers and had similar sample averages for the age of the household head (45.5) and the size of the household (4.7). The mean grower in my sample had about 2 hectares less coffee than the mean grower in Fort and Ruben and a third of a hectare more in other crops, though the means are likely statistically indistinguishable given the large standard deviations in Table 1.

Table 2 summarizes basic information about the three groups (districts). Pruning first started in Group 1 (district 1) in 2005 while the first prune in Groups 2 and 3 was in 2007. Financing from the Peruvian government social development fund FONCODES allowed the NGO implementing the project to coordinate and fund field visits of growers in Group 3 to early adopters in Group 1, which is why the year when Group 3 first had access to a two-year-old prune was in 2008, not in 2009 as the timing of the first prune in the group would suggest.

5. Pruning, yields, and visiting neighbors

A grower who starts systematic pruning in Year 0 could have less production in Year 1 because a plant pruned after the harvest in Year 0 produces nothing in Year 1, though it's possible that unpruned plants nearby produce more because more air and light reach them. The project's default recommendation was for a grower to prune a share of the farm every year, which implies that pruning occurs every year and, assuming a third of the farm is pruned each time, that a plant is pruned once every three years. If a grower follows the recommendation, the yield for the entire farm (Kg/Ha) will increase in Year 2 because plants pruned in Year 0 return to production, though the gains are partially offset by the loss of plants pruned in Year 1. The yield for the farm should be highest in Year 3 as plants pruned in Years 0 and Year 1 will be in production—only those pruned in Year 2 will not produce. Although the yield for the entire farm will be the highest in Year 3, an observer can verify the effects of pruning by looking at the plants pruned in Year 0 two years later, which is why the uptick in adoption should occur then and not in Year 3. Fig. 2 depicts the reported dynamics of

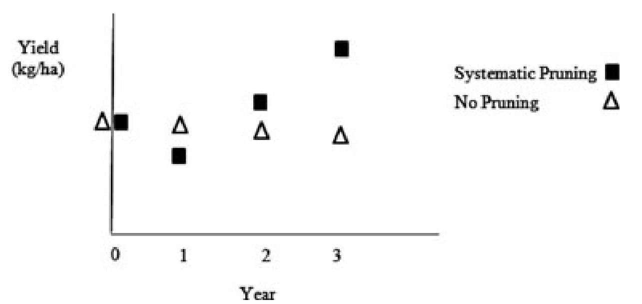

Note: The figure depicts the projected yields for pruning a third of all plants every year beginning after the harvest of Year 0.

Fig. 2. Systematic pruning and yields.

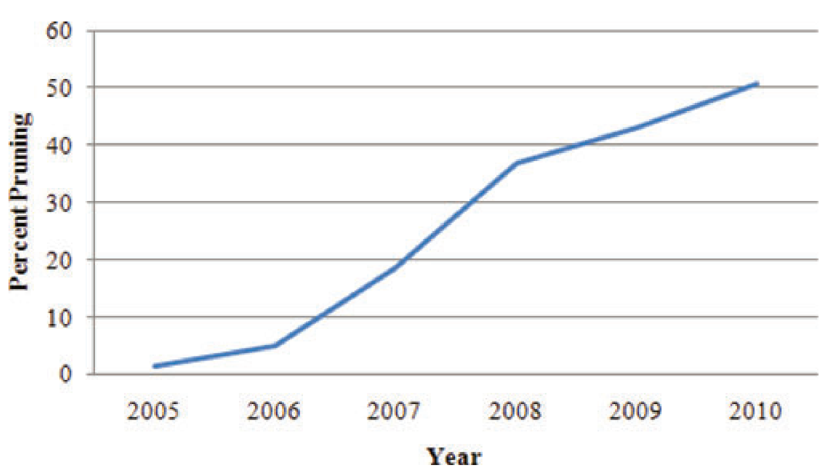

Fig. 1. Cumulative percent of growers pruning ($n = 315$).

systematic pruning supposing that a grower starts pruning after the harvest in Year 0 and prunes a third of the farm each year.

Although Fig. 2 traces the dynamic effects of following the project's recommendation, in practice not every adopting grower pruned every year following the initial adoption. Some growers chose to prune all plants available for pruning in the beginning year, leaving no plants to prune in the following year. This occurred with several growers who had a parcel of very young plants not ready to be pruned (less than five years old) and a parcel with old plants that badly needed pruning. Borrowing constraints provide another reason why not all adopting growers followed the project's recommendation to prune every year. In the conceptual model, a grower's profits in a given year must meet the household's liquidity needs. Changing circumstances may mean that a household that pruned last year forgoes pruning in the current year to maximize production and income in the short term.

Measuring the effect of pruning with nonexperimental data is challenging, especially in the absence of plot-level data. For a perennial crop like coffee, yields in one year reflect activities in previous years like fertilizing, controlling pests, and new plantings. Perhaps more importantly, pruning decisions are endogenous in that growers decide where and how much to prune based on their unique understanding of each plot. Nonetheless, there is clear evidence that growers with sustained participation in the project, who tended to prune, saw yields increase over time relative to growers that exited the project in the first few years. Using a grower-fixed effects model, Barham and Weber (2012) show that by 2010 project participants had yields roughly 50% higher than growers who exited early on. The project's success at improving grower yields helped to make pruning a central part of the Peruvian National Coffee Board's strategy to improve productivity nationwide (Junta Nacional del Café—Peru, 2011)

The 2010 survey asked growers if they knew who were the first growers to prune in their locality (*caserio*), and if so, to name a grower. Then, they were asked if they had visited the parcel of the grower. Of the 236 growers surveyed, 81% said that they had visited an early adopter while 73% gave the name of the adopter. Four growers accounted for almost 40% of the names given. Had the survey asked for more than one name, it is possible that even more growers would have given the name of one of the four early adopters.

The first adopter had a larger than normal coffee farm (over seven hectares of coffee) and a baseline productivity almost double that of the sample median, characteristics that would have made him more likely to adopt and perhaps also more credible in the eyes of visiting growers. One of the other adopters often visited, however, had only two hectares of coffee and started with extremely low yields. He may have adopted because the contribution of coffee to his income had already dropped to dismal levels, so he had little to lose. Furthermore, his success with pruning may have been even more convincing than the success of the earliest adopter who had a higher performing farm even before the project. Though a fascinating area for research, without more information or cross-sectional variation it is impossible to identify the motivation for the earliest adopters to adopt or how their characteristics affected their influence on neighbors. Consequently, the next section focuses on evidence that the initial prune in a group induced the adoption pattern expected given the delayed revelation of information from a pruned plant.

6. Who prunes? An empirical adoption model with neighborhood dynamics

6.1. Model specification

Some learning and technology adoption work (Foster and Rosenzweig, 1995; Cameron, 1999) has studied how planting history, for example, affects the intensity of adoption. In the time-frame of this study growers generally pruned the same area each year as extension agents recommended: of the 52 growers who pruned in 2007, 43 pruned the same amount in 2008. And while the conceptual model allowed for individual learning, once someone in the group has adopted, there is little incentive for a grower to experiment on his own, given that the neighbor's earlier experiment will provide information sooner than his own experiment. The empirics therefore focus on whether or not growers who have not yet pruned begin to prune in response to the maturation of plants pruned by early adopters.

Building on the conceptual model in Section 3, a grower will prune if the expected present value of utility from pruning exceeds that of not pruning. Supposing that a grower has not pruned before, I use the binary variable s_{it} defined in (9) to indicate if a grower starts pruning in period t.

$$s_{it} = \begin{cases} 1 & \text{if } V_{it+1}(a_{it}, .) - V_{it+1}(0, .) > 0 \\ 0 & \text{otherwise} \end{cases}. \tag{9}$$

I rewrite the difference in conditional value functions as a parametric function $D(.)$ which includes as arguments a vector of variables related to $\widetilde{y}_{it}$ presented in Section 3. Let $\widetilde{Y}_{it}$ be a vector of binary variables based on the years since the first prune in the group. A grower enters the panel when at least one grower in the group starts to prune[5], and a grower-year observation falls into one of six exhaustive and mutually exclusive states: the first prune in the group occurs in the current year or the first prune occurred one, two, three, four, or five years ago. Consequently, $\widetilde{Y}_{it}$ includes binary variables for five of the six cases where, for example, $\widetilde{y}_{i0}$ equals one if the first prune in grower *i's* group occurred in the period in question, otherwise it equals zero. The visit by Group 3 to Group 1 is included in the $\widetilde{y}_{it}$ variables by

[5] The only caveat is for Group 1 where the project started working with two communities early enough in 2005 for growers to prune in that year. Growers outside of these communities started working with the project and were exposed to the pruning practice for the 2006 coffee season but not for the 2005 season. To capture this, growers in the two initial communities enter the panel in 2005 while the rest of growers in Group 1 enter in 2006.

supposing that the visiting group has the same information as the group being visited.

I write $D(.)$ as a function of the $\widetilde{y}_{it}$ variables and other relevant variables:

$$V_{it+1}(s_{it},.) - V_{it+1}(0,.) = D\left(\sum_{j=0}^{5} \eta_j \widetilde{y}_{it} + \beta W_i + \alpha K_i + \tau_t + \gamma_i\right). \quad (10)$$

A time effect (τ_t) captures shocks like changes in prices that affect all growers similarly and a grower specific effect (γ_i) allows for unobserved grower hetereogeneity. To capture time invariant characteristics that affect production, I include in K_i the farm's yield (Kg/Ha) in the year prior to entering the project and the farm's elevation (thousands of meters above sea level). In addition to capturing the quality of the plant stock, yields can also reflect grower ability and management intensity (e.g., fertilizer use). Elevation captures agro-climatic conditions important for growing coffee.

In W_i I include measures of household wealth, portfolio diversification, labor supply and the life cycle stage: *Total Land*, *Gross Income—Other Crops*, *Labor,* and a linear and quadratic term for *Age* (of the household head).[6] *Labor* is the number of persons in the household ages 10 to 65—in essence, everyone who could help with farm work. All variables other than the $\widetilde{y}_{it}$ variables are from the baseline survey and are time invariant.

Having data on all growers who initially participated in the project allows me to use the actual behavior of nearby growers instead of what a grower reports about his neighbors. In a study of network data from Kenya, Hogset and Barrett (2009) found that farmers' beliefs about their neighbors' decisions were not reliable, where reliability is the correlation between what a farmer thinks his neighbor did and what his neighbor reported doing. More importantly, a farmer's error in predicting what his neighbor did was correlated with the farmer's behavior; a farmer tends to report that his neighbors made decisions similar to his. I avoid this problem by using a neighbor-related variable that only depends on the timing of the first prune in the group, which comes from the project's pruning registry.

6.2. Identification

A challenge in taking Eq. (10) to the data is that the random variable μ_{it}, which reflects expectations about the effects of pruning, may respond to an unobservable group shock $\widetilde{u}_{it}$ in addition to the $\widetilde{y}_{it}$ variables, which only vary at the group level. For example, growers in the same group may prune more because they received more guidance from project extension agents. To separate the effect of the timing of the first prune in the group from a group effect, I rewrite Eq. (5) from Section 3 as

$$\mu_{it} = I(\widetilde{y}_{it}) + \widetilde{u}_{it}. \quad (11)$$

I define the group shock as a function of the group adoption rate and an idiosyncratic error term.

$$\widetilde{u}_{it} = \theta_1 \overline{g}_{-i} + \varepsilon_{it}, \quad (12)$$

where the average group adoption rate is defined as $\overline{g}_{-i} = \sum_{t=\text{first prune}}^{T} g_{-it}$, with g_{-it} being the adoption rate for group g at time t (excluding the decision of grower i) and "t = first prune" is the year when pruning started in the group.

Specifying an unobservable variable ($\widetilde{u}_{it}$) as a function of time averages of observable variables is a common approach in panel data models. Papke and Wooldridge (2008) address endogeneity in panel models by specifying the random effect as a function of time averages of potentially endogenous variables, a method known as the Mundlak-Chamberlain device (Mundlak, 1978; Chamberlain, 1982). Lewis et al. (2011) use an approach similar to the one used here to study spatial spillovers in the expansion of organic dairy farming.

In his 1993 seminal paper, Manski discusses the challenge of identifying endogenous group effects (an individual's behavior moves with the group average) from contextual effects (an individual's behavior varies with exogenous characteristics of the group) and correlated effects (all individuals behave similarly because they face a similar environment). Including the average group adoption rate controls for time invariant group characteristics that affect pruning decisions while year fixed effect terms control for shocks that affect all groups similarly.

Conley and Udry (2010) also control for the effect of similar environmental conditions on the decisions of farmers in the same proximity. Technologies vary in their sensitivity to local conditions. Fertilizer use, which Conley and Udry study, is a sensitive outcome to study because the marginal productivity of fertilizer, and therefore the optimal use, can vary substantially depending on weather and pest conditions. Although environmental conditions clearly affect yields, the agronomic recommendation to prune older plants is not conditional on rainfall, pests, or other shocks. Nonetheless, there is still a concern that the error in Eq. (12) is correlated with the $\widetilde{y}_{it}$ variables. The advantage of the present case is that the maturation of pruned trees implies a distinct adoption pattern if social learning is important. Aside from the signal sent by a neighbor's actions, nothing about the effect of pruning is learned from viewing a recently pruned plant; an effect of $\widetilde{y}_{i0}$ on pruning possibly captures a mixture of imitation, the presence of pioneers, and group specific shocks. Little is learned from viewing a one-year old prune, so $\widetilde{y}_{i0}$ should have little effect on the adoption decisions of growers who have not already adopted. Because plants fully recover in the second year, access to a two-year old prune, captured by $\widetilde{y}_{i2}$, should encourage adoption, especially relative to access to a one-year-old prune.

[6] Education may also affect adoption. Years of schooling completed by the household head is available for the 236 growers in the 2010 survey. Average schooling for these growers is 7.1 years with a standard error of .28, implying little variation in schooling. Furthermore, education is highly correlated with age (younger household heads have more education) so the model's age terms will capture much of the variation in education.

Finding a strong increase in adoption in Year 2 relative to Year 1 identifies social learning if there is no unobservable shock that follows a similar pattern. To imitate social learning, correlated unobservable variables would have to motivate some growers to prune initially, have little effect in the following year, and then reappear to motivate more pruning in the year after that. Identifying social learning in this case, therefore, only requires that unobservable shocks cause a different adoption pattern than the one expected given the *a priori* knowledge about the information revealed by a pruned plant.

Because the decision modeled is that of pruning for the first time, a grower leaves the sample once having pruned. One may ask how conditioning on not having pruned previously affects the interpretation of estimation results. Brock and Durlauf (2010) discuss the observable effects of social influences under various assumptions. One of their key assumptions is that lower ability agents adopt later when a greater percentage of the group has already adopted. Applying the assumption to the present case, suppose that higher ability growers are more likely to adopt at any given time because they can earn higher profits from pruning. As higher ability growers adopt and leave the sample, the remaining low ability growers would be less responsive to new information, which means that the effects of the maturation of neighbors' prunes on growers who have not yet adopted would be weaker than for the entire sample. Alternatively, lower-ability growers may have less private information and be more responsive to information from peers.

6.3. Estimation

Assuming that ε_{it} is an i.i.d. normally distributed random variable with a mean zero and standard deviation σ_ε, the probability of pruning for the first time is

$$\Pr\left(s_{it} = 1 \,|\, \widetilde{y}_{it}, W_i, K_i, \overline{g}_{-i}, \tau_t, \gamma_i, \sum_{\tau=0}^{t-1} a_{i\tau} = 0\right) = \Phi\left(\sum_{j=0}^{5} \eta_j \widetilde{y}_{ij} + \beta W_i + \alpha K_i + \theta_1 \overline{g}_{-i} + \tau_t + \gamma_i\right), \quad (13)$$

where $\Phi(.)$ is the standard normal cumulative distribution function. I treat the grower specific effect γ_i as an i.i.d. normal random effect with mean zero and standard deviation σ_γ. Modeling grower heterogeneity as a random effect allows growers who adopt at the first opportunity and growers who never adopt to contribute to identification of parameter estimates. As a robustness check, I later treat γ_i as a fixed effect in a linear model.

Equation (13) is then estimated as a Random Effects Probit model using Maximum Likelihood and the Gauss-Hermite quadrature method. I use a likelihood ratio test for the restriction that σ_γ equals zero, which I fail to reject at the 10% confidence level. I therefore exclude the grower specific random effect in estimation.

The average group adoption rate ($\overline{g}_{-i}$) captures correlation in decisions between growers in the same group. To be clear, a nonadopter and an adopter in the same group will have different, though highly correlated values for $\overline{g}_{-i}$. Because the $\widetilde{y}_{it}$ variables exhaust all possible cases, the variable corresponding to Year 1 ($\widetilde{y}_{i1}$) is omitted. For all results, I calculate robust standard errors clustered by grower to account for correlation in errors across years. Because information is assumed to be correlated among growers in the same district, it is also sensible to cluster standard errors at the district level; doing so gives the same qualitative results.

I check the estimates from the base model (model 1) in four ways. The first three checks (models 2–4) assume the form of social learning in the base model—that what matters for adoption is when the grower gains access to a two-year old prune of an early adopter. The first check (model 2) tests if the results are sensitive to the assumption about ε_{it} by assuming that it has a logistic instead of a normal distribution. The second check drops the grower-year observations where information on pruning does not exist, and it was assumed that the grower did not prune. In the third check (model 4) I estimate a linear probability model with grower fixed effects.

In contrast to the first three checks, the fourth check (model 5) assumes a social learning model where the accumulated experience of others in a grower's group matters, not the timing of the first prune in the group. Foster and Rosenzweig (1995) illustrate this common approach to social learning; to identify the effect of learning from others, they use the average lagged cumulative sum of hectares cultivated under high yielding seed varieties (the technology in question) of farmers in the village. Taking a similar approach, I replace the discrete year variables indicating the time since initial adoption with the lagged cumulative adoption rate of the group. If what matters for adoption is access to a two-year old prune, then the lagged accumulated group experience will likely be negatively correlated with adoption. After the initial wave of adoption in Year 0, few growers adopt in Year 1 since there is little additional information revealed, causing a negative relationship between adoption in Year 0 and new adoption in Year 1. Furthermore, there will likely be a negative relationship between cumulative adoption in Year 1, which will be low because of little addition adoption in that year, and adoption in Year 2, which will be high if access to a two-year old prune is important.

7. Results and discussion

Tables 3 and 4 presents the marginal effects calculated at the means for continuous variables and for a discrete change in the binary variables in $\widetilde{Y}_{it}$.

7.1. Temporal patterns in adoption

The marginal effects for the $\widetilde{y}_{it}$ variables in model 1 suggest that growers are marginally less likely to adopt in Year 1, the excluded year, than in Year 0, defined as the year when someone in the group starts to prune. This is consistent with the

Table 3
Marginal effects for models 1 through 4 (Dependent variable: indicator if pruned)

Variable	Base (1)	Logistic (2)	Dropping missing Obs. (3)	Linear fixed effects (4)
Year 0	0.084	0.064	0.124*	0.13***
	(0.064)	(0.055)	(0.070)	(0.05)
Year 2	0.182***	0.150***	0.229***	0.37***
	(0.035)	(0.035)	(0.040)	(0.04)
Year 3	0.003	0.004	−0.001	0.20***
	(0.044)	(0.037)	(0.048)	(0.03)
Year 4	0.159***	0.141***	0.149**	0.33***
	(0.055)	(0.049)	(0.063)	(0.05)
Year 5	0.247***	0.204***	0.212**	0.44***
	(0.076)	(0.065)	(0.087)	(0.06)
Calendar Year 2005	−0.071	−0.089	−0.198*	−0.15***
	(0.095)	(0.089)	(0.102)	(0.05)
Calendar Year 2006	0.142*	0.113*	0.068	0.18***
	(0.076)	(0.064)	(0.086)	(0.05)
Calendar Year 2007	0.092	0.074	−0.009	−0.00
	(0.065)	(0.057)	(0.076)	(0.04)
Calendar Year 2008	0.177***	0.145***	0.108	0.12***
	(0.061)	(0.054)	(0.073)	(0.03)
Calendar year 2009	0.058	0.038	0.023	0.03
	(0.051)	(0.047)	(0.061)	(0.03)
Average Group Adoption	0.114	0.005	0.270	
	(0.153)	(0.148)	(0.171)	
Total land	0.002*	0.002*	0.003**	
	(0.001)	(0.001)	(0.001)	
Gross income—other crops (/1000)	0.001*	0.001**	0.001	
	(0.001)	(0.001)	(0.001)	
Yield (/1000)	0.018***	0.014***	0.020**	
	(0.007)	(0.005)	(0.008)	
Elevation (/1000)	0.088	0.081	0.103	
	(0.067)	(0.055)	(0.078)	
Labor	0.000	−0.002	0.002	
	(0.004)	(0.003)	(0.005)	
Age	−0.002	−0.001	−0.001	
	(0.004)	(0.003)	(0.004)	
Age squared (/100)	0.002	0.001	0.001	
	(0.004)	(0.003)	(0.005)	
Observations	1,300	1,300	1,134	1,336
Growers	307	307	307	315
Pseudo R^2	0.194	0.197	0.237	0.228[a]
P-value fo Ho: Year 0 = Year 2	0.084	0.078	0.095	0.000

***$P < 0.01$, **$P < 0.05$, *$P < 0.1$ Robust covariance matrix clustered at the grower level is calculated. Standard errors for the marginal effects (in parenthesis) are then calculated using the delta method.
[a]Adjusted R^2.

understanding that little can be learned by observing a plant pruned a year ago. But the next year (Year 2), when growers have access to a two-year-old prune, adoption expands. Models 2 and 3 give results similar to model 1: there is more adoption in Year 0 than in Year 1. But the next year, when growers have access to a two-year old prune, there is an expansion of adoption that is not observed from Year 2 to Year 3. The linear fixed effects model (model 4), which does not require distributional assumptions about ε_{it} or a lack of correlation between the covariates and the grower specific term to estimate consistent coefficients, yields results qualitatively similar to models 1 through 3. The marginal effects from models 1 through 4 imply that a grower has a 0.15 to 0.37 higher probability of starting to prune in Year 2 relative to Year 1. The estimated learning effect is similar in size to Conley and Udry's (2010) estimate of the effect of observing neighbors having lower than expected profits: they find that a one standard deviation in a farmer's observation of bad news (from neighbors) at a previously used fertilizer level increased the probability that the farmer would change his fertilizer use by 0.15.

Table 4
Marginal effects for model 5 – an alternative model of social learning (Dependent variable: indicator if pruned)

Variable	Model 5
Lagged Cumulative Adoption Rate	−0.625***
	(0.113)
Calendar Year 2005	−0.477***
	(0.058)
Calendar Year 2006	−0.329***
	(0.057)
Calendar Year 2007	−0.162***
	(0.047)
Calendar Year 2008	−0.134***
	(0.039)
Calendar Year 2009	−0.075***
	(0.029)
Average Group Adoption	0.406***
	(0.138)
Total Land	0.001
	(0.001)
Gross Income from Other Crops (/1000)	0.003
	(0.002)
Yield (/1000)	0.032**
	(0.013)
Elevation (/1000)	0.064
	(0.073)
Labor	0.003
	(0.005)
Age	−0.003
	(0.004)
Age squared (/100)	0.005
	(0.004)
Observations	1,188
Growers	301
Pseudo R^2	0.181

***$P < 0.01$, **$P < 0.05$, *$P < 0.1$. Robust covariance matrix clustered at the grower level is calculated. Standard errors for the marginal effects (in parenthesis) are then calculated using the delta method.

Compared to models 1 and 2, models 3 and 4 show stronger evidence that adoption is more likely in the year of the first pruning in the group than in the following year, though they also show a greater coefficient on Year 2 such that the difference between coefficients for Year 0 and Year 2 are similar in model 1 and model 3. In all four models, testing for equality of coefficients reveals that at the 10% level the effect in Year 2 is statistically different from and greater than the effect in Year 0.

For correlated unobservable variables to cause the adoption pattern observed in the first three years, a shock must induce more growers to prune initially, disappear in the intervening year and then reappear in the following year. Such a shock seems unlikely, especially because the unique revelation of information from pruning provides a strong *a priori* reason to expect the jump in adoption from Year 1 to Year 2 found in all models. Furthermore, the jump in adoption in Year 2 occurs even after controlling for the calendar year, which captures price and weather shocks that could cause adoption to rise or fall in a given year.

The estimated coefficients on Years 3, 4, and 5 imply that adoption slows in Year 3 but then may increase again in Years 4 and 5. As pruned plants from earlier waves of adoptions mature, growers who still have not pruned would have more plants to observe and would have to travel less to observe them. However, it is best to take estimates for the later years with caution since they are based on fewer observations than earlier years: for Year 3, 4, and 5 the number of observations is 207, 180, and 120.[7] Furthermore, bias from poor recall data in 2009 and possibly better reporting for the 2010 year will primarily affect later years (recall that 93% of growers had already had access to a two-year old prune by 2009 when only recall data are used).

A pertinent question is whether the data also support an alternative model of social learning where the group's accumulated experience in pruning matters most for adoption. I replace the discrete year variables indicating the time since initial adoption with the lagged cumulative adoption rate of the group. A positive relationship between cumulative group adoption and an individual's behavior is often interpreted as evidence of social learning. In the previous section I noted why the lagged recovery of a pruned plant could induce a negative correlation between cumulative group adoption and an individual grower's adoption. This appears to be the case, with the marginal effect of the lagged cumulative adoption rate being negative, quite large, and statistically different from zero (Table 4).

7.2. *Wealth and liquidity*

The results for the control variables from the baseline suggest that wealth, liquidity, and perhaps ability affect adoption, though the estimated marginal effects are small. Increasing total land by a hectare increases the probability of adopting in any given year by 0.002. *Gross Income – Other Crops* is significant in models 1 and 2, which is consistent with comments by project extension agents stressing the value to small-scale growers of a consistent cash flow from selling bananas and avocados, and that growers feared that pruning would initially depress production. The effect, however, is modest, with the largest estimate implying that each thousand dollars in extra gross income from other crops increases the probability of pruning by 0.001.

Because the liquidity variables reflect conditions before the project, there can be no reverse causality (e.g., growers who expect to prune plant banana and avocado trees). That said, the variables may be correlated with unobservable characteristics. For example, risk-averse growers may have more banana and avocado trees and therefore higher gross income from other crops.

Growers with a higher yield at the beginning of the project were more likely to prune—according to model 1, increasing yields by 50% of the median baseline yield increases the probability of pruning by 0.0045. The positive correlation indicates that productivity reflects more than the opportunity cost of pruning. It may also capture liquidity constraints since high yields from previous years could mean more cash for the present year. But like other wealth variables, initial yields are likely correlated with unobservable characteristics such as management ability or intensity. One would expect growers who more intensely managed their coffee plants before the project to adopt practices that further increase intensity.

8. Conclusion

Coffee growing has in some times and places in Latin America fostered and sustained a rural middle class—a noteworthy association in a region plagued by poverty and inequality. Recent research on the incomes that small holders earn from coffee show that the effect of yields (coffee produced per hectare) on income can dominate that of higher prices earned by selling to certified markets like Fair Trade and organic (Barham et al., 2011; Barham and Weber, 2012).

This study shows how a major source of variation in productivity, the systematic pruning of coffee plants, spread among growers in central Peru. The strong increase in adoption when the first pruned plants in a group recovered provides evidence for the social nature of technology diffusion. The findings also show differences in household willingness to incur the short run costs associated with long run yield and income gains. With incomplete rural credit markets, households with small and undiversified portfolios may find themselves earning less and less as they postpone pruning or replanting. Being surrounded by other low-income growers who passively manage their farms (no pruning, for example) can decrease the likelihood that any one grower adopts a different, more intensive management system. And if no one innovates, growers cannot learn from each other to improve yields and incomes.

Despite the evidence that seeing is believing, a caution on the use of demonstration plots in agricultural extension programs is needed. Growers in the area tended to view demonstration plots with skepticism, doubting whether they can achieve similar results on their own. Having a neighbor who successfully implements a practice is more convincing. The empirical results suggest that the most expedient approach to speed diffusion may be to reward several centrally located members for experimenting with a method. Cash incentives, however, may be insufficient to ensure appropriate and continued application of a technology. In such cases, gaining credibility with potential early adopters may play a more important role, which was

[7] The decrease in grower-year observations is from adopters dropping out and because not all the groups reached years four and five in the study period.

clearly the case for the project studied, which did not use cash incentives but instead established a consistent presence in the communities of participating growers. This suggests that the social capital of organizations diffusing technology may matter a lot.

Although growers may have to see actual results before adopting, second and third-party diffusion of information through conversation, the radio, or extension bulletins like the one produced by the Peruvian National Coffee Board (Junta Nacional del Café—Peru, 2011) raise awareness about new practices even though they may be insufficient to convince growers of their effectiveness. Growers may then want to visit early adopters to see the results for themselves as has happened in the most accessible community of the project, which has received visits from extension agents and coffee growers from around Peru. Thus, one can easily imagine how traditional forms of acquiring knowledge such as extension workshops complement social learning by motivating farmers to visit peers.

Policy makers have limited or no control over differences in yields that stem from prices, land endowments, and weather. Governments and NGOs, however, can affect how rural households produce. Given the reported effect of pruning on yields, it is surprising that the practice did not previously enter Peru from Costa Rica or Colombia where versions of it are common practice. It took two private companies paying agricultural extension agents to start the diffusion of systematic pruning in central Peru. Now the practice has spread to growers in northern Peru who work with a branch of the NGO that implemented the project in central Peru and, as mentioned earlier, it forms part of a national strategy to renovate plantations and improve yields. One of the higher return policy interventions for alleviating rural poverty may be fomenting innovation in the sectors where the poor derive their livelihoods. The case of systematic pruning in central Peru provides an example where private actors successfully undertook such an intervention.

Acknowledgements

I thank Brad Barham, Jean-Paul Chavas, Laura Schechter, Hope Michelson, seminar participants at the University of Wisconsin-Madison and the Midwest International Economic Development Conference, and three anonymous reviewers for helpful comments. The work behind the study was made possible by cooperation by Lavazza, Volcafe, Alvaro Llobet, Orlando Diaz, Javier Salazar, Diana Bazan Vargas, and Antero Flores Nano and financial support from the Graduate School of UW-Madison, the Peruvian National Coffee Board (Junta Nacional del Café), and Rainforest Alliance. All errors are my responsibility.

References

Agrobanco, 2011. Crecimiento Sostenido en Plantaciones de Café Debido a la Aplicación de la Poda Sistemática. Available at http://www.agrobanco.com.pe/news/files/175.pdf, Accessed June 2012.

Bandiera, O., Rasul, I., 2006. Social networks and technology adoption in northern Mozambique. Econ. J. 116, 869–902.

Barham, B.L., Callenes, M., Lewis, J.M., Gitter, S., Weber, J., 2011. Fair Trade/Organic Coffee, Rural Livelihoods, and the 'Agrarian Question': Southern Mexican Coffee Families in Transition. World Dev. 39(1), 134–145.

Barham, B.L., Weber, J.G., 2012. The economic sustainability of certified coffee: Recent evidence from Mexico and Peru. World Dev. 40(6), 1269–1279.

Besley, T., Case, A., 1994. Diffusion as a Learning Process: Evidence from HYV Cotton. Discussion Paper 174. Available at http://www.princeton.edu/rpds/papers/pdfs/besley_case_diffusion.pdf, Accessed June 2012.

Brock, W.A., Durlauf, S.N., 2010. Adoption curves and social interactions. J. Eur. Econ. Assoc. 8, 232–251.

Cameron, L.A., 1999. The importance of learning in the adoption of high-yielding variety seeds. Am. J. Agric. Econ. 81, 83–94.

Chamberlain, G., 1982. Multivariate regression models for panel data. J. Econometrics. 18, 5–46.

Chamely, C.P, 2004. Rational Herds: Economic Models of Social Learning. Cambridge University Press, Cambridge.

Conley, T.G., Udry, C.R., 2010. Learning about a new technology: Pineapple in Ghana. Am. Econ. Rev. 100, 35–69.

Fort, R., Ruben, R., 2009. The impact of Fair Trade certification on Coffee producers in Peru. In: Ruben, R. (Ed.), The Impact of Fair Trade. Wageningen Academic Publishers, Wageningen.

Foster, A.D., Rosenzweig, M.R., 1995. Learning by doing and learning from others: Human capital and technical change in agriculture. J. Pol. Econ. 103, 1176–209.

Hogset, H., Barrett, C.B., 2009. Social learning, social influence and projection bias: A caution on inferences based on proxy–reporting of peer behavior. Econ. Dev. Cult. Change 58, 563–589.

Junta Nacional del Café—Peru, 2011. Propuesta de Incremento de Productividad para la Caficultura Peruana. Available at http://progresonetwork.ning.com/profiles/blogs/propuesta-de-incremento-de#comments, Accessed June 2012.

Lewis, D.J., Barham, B.L., Robinson, B., 2011. Are there spatial spillovers in the adoption of clean technology? The case of organic dairy farming. Land Econ. 87(2), 250–267.

Maertens, A., 2009. Who Cares What Others Think (or Do)? Social Learning, Social Pressures and Imitation in Cotton Farming in India. Job Market Paper. Applied Economics and Management, Cornell University.

Manski, C.F., 1993. Identification of endogenous social effects: The reflection problem. Rev. Econ. Stud. 60, 531–542.

Miller, N.F., Dollard, J., 1941. Social learning and imitation. Yale University Press, New Haven.

Moser, C.M., Barrett, C.B., 2006. The complex dynamics of smallholder technology adoption: The case of SRI in Madagascar. Agric. Econ. 35, 373–88.

Mundlak, Y., 1978. On the pooling of time series and cross section data. Econometrica 46, 69–85.

Munshi, K., 2004. Social learning in a heterogeneous population: Technology diffusion in the Indian green revolution. J. Devel. Econ. 73, 185–213.

Papke, L.E., Wooldridge, J.M., 2008. Panel data methods for fractional response variables with an application to test pass rates. J. Econometrics. 145, 121–33.